Benchmarking pthreads

Bronis R. de Supinski and John May

Center for Applied Scientific Computing

March 18, 1999

Measuring thread performance

- Threading options
  - pthreads
  - OpenMP
  - Threaded libraries
- How to choose
  - Ease of use
  - Performance
- No threads benchmark suites!!

Using SKaMPI

- Advantages
  - Extensible interface
  - Data collection and test management facilities
  - Will support mixed models benchmarks
- Drawbacks
  - Times single measurement action
  - Relies on MPI
  - Design flaws
- Challenges
  - Termination
  - CPU bindings

Thread creation

- Measuring thread creation cost is difficult
  - Number of threads is limited
  - Overhead
  - Solution: recursive thread creation
- Using default attributes: 118.7 μs
- Eventually will vary:
  - Detached state
  - Contention scope
  - Stack size?

Context switch

- **Bind threads to same CPU**
- **Repeatedly call sched_yield**
- **Running thread alternates quickly**
- **Measured time: 4.4 μs**

Thread communication

- LogP model
  - Latency (wire time)
  - Overheads
  - Gap

- Measure
  - Ping-pong time: $2*(o_S + L + o_R)$
  - Repeatedly “sending”: o_S
  - Repeatedly “receiving”: o_R
  - Calculate L
  - Applicability of gap?

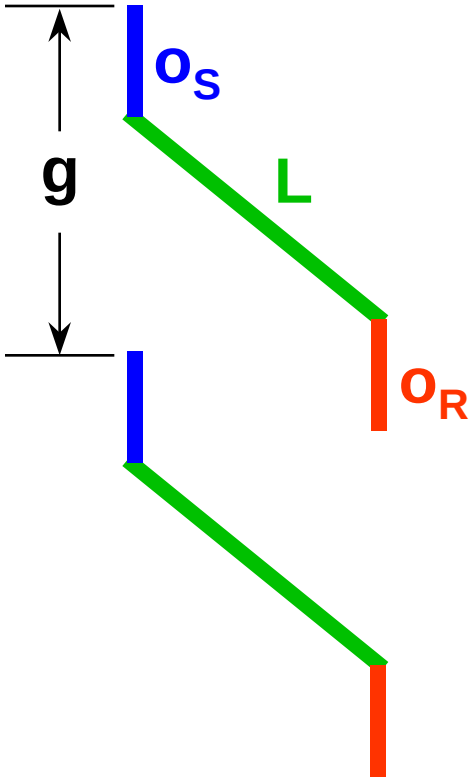

Conditions

- Operations
  - pthread_cond_signal
  - pthread_cond_wait
- Associated mutex
- Ping-pong measurements
  - Unbound: 48.9 µs (Sun: 25.0 µs)
  - Same CPU: 29.2 µs (Sun: 25.7 µs)
  - Different CPUs: 74.3 µs (Sun: 25.2 µs)
- Overheads
  - Signal: 0.606 µs
  - Wait: Different CPUs: 45.7 µs

Mutexes

- Ping-pong obstacles
  - Non-determinism
  - Idempotency required
- Measurements
  - Unbound: 3.7 µs
  - Same CPU: 37.8 µs
  - Different CPUs: 3.7 µs
  - No contention: 0.638 µs

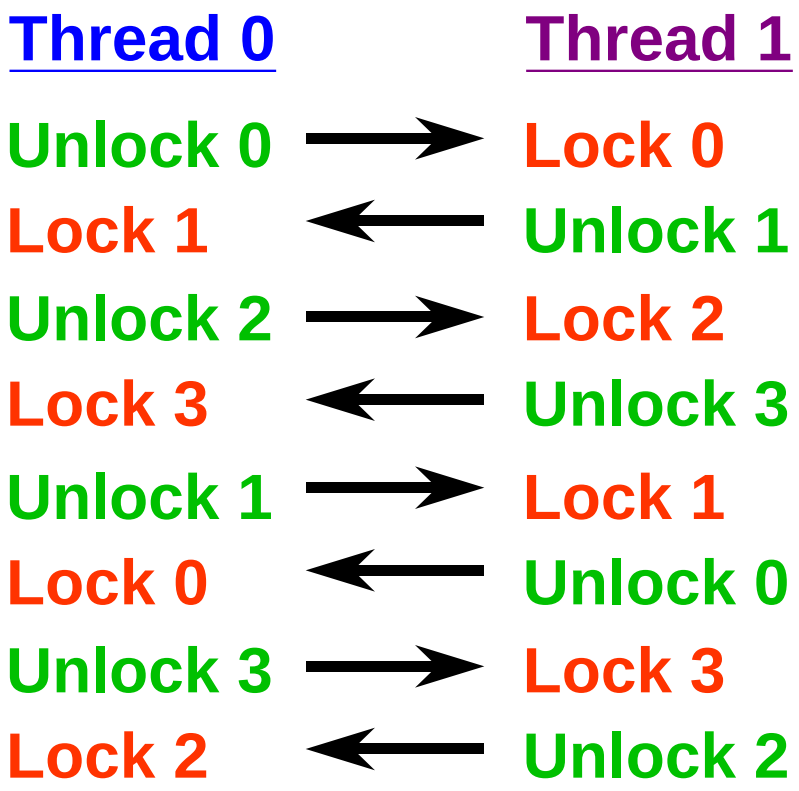

Summary

- SKaMPI modifications
  - Fills threads microbenchmark void
  - Can extend to mixed models
- Basic aspects of pthreads performance
  - Creation
  - Context switch
  - Timeslice
  - Communication
- IBM's implementations
  - Mutexes are good
  - Conditions could be improved
- Plan to add other tests, variations

UCRL-MI-133178

Work performed under the auspices of the U. S. Department of Energy by Lawrence Livermore National Laboratory under Contract W-7405-Eng-48